ARI Research Note 97-20

A Cognitive Architecture for Solving Ill-Structured Problems: Final Report

Keith J. Holyoak
University of California at Los Angeles

Paul Thagard
Princeton University

Research and Advanced Concepts Office
Michael Drillings, Chief

August 1997

United States Army
Research Institute for the Behavioral and Social Sciences

U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency Under the Jurisdiction of the Deputy Chief of Staff for Personnel

EDGAR M. JOHNSON
Director

Research accomplished under contract
for the Department of the Army

University of California at Los Angeles

Technical review by

Michael Drillings

NOTICES

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED (from. . . to)
1997, August	Final	September 1986-August 1987

4. TITLE AND SUBTITLE A Cognitive Architecture for Solving Ill-Structured Problems: Final Report	5a. CONTRACT OR GRANT NUMBER MDA903-86-K-0297 5b. PROGRAM ELEMENT NUMBER 0601102A
6. AUTHOR(S) Keith J. Holyoak (UCLA) and Paul Thagard (Princeton University)	5c. PROJECT NUMBER B74F 5d. TASK NUMBER 8513 5e. WORK UNIT NUMBER C71
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Psychology University of California at Los Angeles Los Angeles, CA 90024	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Institute for the Behavioral and Social Sciences ATTN: PERI-BR 5001 Eisenhower Avenue Alexandria, VA 22333-5600	10. MONITOR ACRONYM ARI 11. MONITOR REPORT NUMBER Research Note 97-20

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

COR: Judith Orasanu

14. ABSTRACT *(Maximum 200 words)*:

A computational theory of analogical mapping is described, based on a set of constraints. The theory is embodied in a computer simulation that is applied to several examples, including psychological data on the mapping process.

15. SUBJECT TERMS

Analogy   Constraint satisfaction   Parallel processing

SECURITY CLASSIFICATION OF			19. LIMITATION OF ABSTRACT	20. NUMBER OF PAGES	21. RESPONSIBLE PERSON (Name and Telephone Number)
16. REPORT	17. ABSTRACT	18. THIS PAGE			
Unclassified	Unclassified	Unclassified	Unlimited	16	

A COGNITIVE ARCHITECTURE FOR SOLVING ILL-STRUCTURED PROBLEMS:
FINAL REPORT

CONTENTS

Page

1. Introduction 1
1.1 Relevance to the ARI Mission 1
1.2 Components of Analogy Use 1
1.3 Accomplishments of the Project 1

2. A Constraint-Based Architecture for Analogy 2
2.1 Parallel constraint Satisfaction 3
2.2 Constraints on analogical Mapping 3
2.3 Analogical Mapping in ACME 5
2.4 Analog Retrieval in ARCS 5

3. Implications for the Army 9
3.1 Training Implications 9
3.2 Potential Aid for Computer-Assisted Problem Solving 10

4. References 11

A Cognitive Architecture for Solving Ill-Structured Problems: Final Report

1 INTRODUCTION

This project was directed at the development of components of a cognitive architecture for modeling how people solve ill-structured problems, ones which cannot be solved by application of routine procedures. The primary focus was on the use of analogies in coordination with rule-based problem-solving strategies. The core of a cognitive architecture consists of three subsystems: a problem-solving system, capable of drawing inferences to construct plans for attaining goals; a memory system, which can be searched in an efficient manner to identify information relevant to the current problem; and an inductive system, which generates new knowledge structures to be stored in memory so as to increase the subsequent effectiveness of the problem-solving system (Holland, Holyoak, Nisbett, & Thagard, 1986). The project investigated all of these component subsystems. As our work has been reported in detail in several papers already submitted to the ARI, the present report will provide a relatively brief overview.

1.1 Relevance to the ARI Mission

Our research has combined the development of explicit computational models of cognitive architecture with allied experimental tests of fundamental theoretical ideas embodied in the models. This research is directly related to the mission of the ARI, since better models of problem solving and learning should lead to strategies for improving the effectiveness of Army personnel. Moreover, our emphasis has been on solving ill-structured problems, of the sort that personnel are likely to face in real situations.

1.2 Components of Analogy Use

Previous research suggests that one of the keys to solving novel problems is to use past solutions to analogous problems to suggest how the new problem might be decomposed (Gick & Holyoak, 1980, 1983; Ross, 1987). Analogical problem solving is a skill that clearly depends on the coordination of problem solving, memory access, and induction. It is useful to decompose the process of analogical problem solving into four steps: *retrieval, mapping, transfer*, and *learning*. The retrieval step involves accessing a plausibly useful source analog in memory. It is particularly difficult to identify candidate source analogs when they are concealed within a large memory base, and when the source and the target problem were encountered in different contexts and have salient dissimilarities. These theoretical issues are closely related to those raised in Schank's (1982) discussion of "reminding". The mapping step requires finding an optimal set of correspondences between the elements of the source and target. The transfer step involves using the mapping to derive useful inferences about the target (e.g., a new subgoal), which can then be tested. Finally, learning can result in the induction of new knowledge structures that summarize the useful commonalities between the source and target that have been discovered, a process we will refer to as *schema induction*.

1.3 Accomplishments of the Project

During the three years of this ARI contract, we made major strides in developing and testing basic components of our proposed architecture. The following accomplishments were especially significant:

(1) We wrote an extended description of the cognitive architecture embodied in the Common LISP program PI (Thagard, 1988a).

(2) We extended the PI cognitive model to include analogical problem solving (Holyoak & Thagard, 1989a). This was the first explicit model to deal with all four major steps in analogical problem solving, and to suggest an integration of analogical and rule-based solution methods.

(3) We reviewed existing research on computational models of analogy (Thagard, 1988b).

(4) We re-evaluated the model of analogy in PI, deciding that its account of analogical mapping was insufficiently constrained to adequately capture some important aspects of human analogical problem solving. Based on an analysis of these constraints (see Section 2.2 below), we then developed a constraint-satisfaction theory of analogical mapping that has been implemented in the program ACME (Analogical Constraint Mapping Engine). This mapping model, which automatically generates a localist connectionist network of mapping hypotheses relating two symbolic structures, was then applied to more than 20 examples from different domains, including simulations of the results of several psychological experiments (Holyoak & Thagard, 1989b).

(5) We applied ACME to a different domain: analogies used in chemical education (Thagard, Cohen, & Holyoak, 1989). This set of applications has direct implications for the analysis of instructional materials for science education.

(6) We extended the principles underlying ACME to generate a new model of analog retrieval, implemented in the program ARCS (Analog Retrieval by Constraint Satisfaction). In essence, ARCS compares in parallel the target to structures stored in memory, and selects a small subset of stored structures that best satisfy the constraints that determine the adequacy of an analogy. The psychological and computational adequacy of ARCS has been tested on four data bases (Thagard, Holyoak, Nelson, & Gochfeld, 1989). Two of these data bases involve direct simulation of the results of psychological experiments, and two test the capacity of the model to deal with retrieval from large data bases.

(7) We wrote an overview of our new models of mapping and retrieval (Holyoak & Thagard, 1990).

(8) We conducted an extensive series of experiments to identify conditions under which analogical transfer can be expected when a substantial delay and context change intervenes between presentation of source analogs and of the target problem (Catrambone & Holyoak, 1988, 1989). The results have direct implications for developing robust training procedures that foster induction of relatively content-independent problem schemas.

(9) We extended the constraint-satisfaction principles embodied in our models of analogical mapping and retrieval to develop a related model of the evaluation of explanations. In the ECHO program (Explanatory Coherence by Harmany Optimization), sets of propositions describing hypotheses and observations, together with explanatory relations among them, provide the input to a set of routines that construct a constraint network. This network embodies basic constraints that govern explanatory coherence (e.g., similarity, generality, and analogy to other accepted explanations). After the network is constructed, it is settled to weed out less coherent explanations of the data. The ECHO model was applied to numerous cases involving scientific and legal reasoning (Thagard, 1989; Thagard & Nowak, 1988). The model potentially offers a highly general tool for modeling decision making, especially in circumstances when detailed quantitative information about probabilities is lacking.

2 A CONSTRAINT-BASED ARCHITECTURE FOR ANALOGY

We will now briefly review the operation of the analogy models we have developed.

2.1 Parallel Constraint Satisfaction

Both Gestalt theorists (Maier, 1930) and current cognitive scientists (Hofstadter, 1984) have emphasized that high-level reasoning has important commonalities with perceptual processing. We believe these commonalities can be explicitly formulated in models that employ a central feature of current connectionist models: *parallel constraint satisfaction*. In general, parallel constraint satisfaction is preferable to any serial decision procedure when: (a) a global decision is composed of a number of constituent decisions, (b) each constituent decision should be guided by multiple constraints, (c) the outcome of the global decision could vary depending on the order in which constraints are applied and constituent decisions are made, and (d) there is no principled justification for preferring any particular ordering of constraints or of constituent decisions. In essence, all of the constituent decisions are made simultaneously and incrementally, with continuous communication of partial results. (See Thagard, 1986, for a philosophical discussion of the importance of parallel computation.)

The process of retrieving analogs stored in memory, and the process of mapping a target analog with a source analog, exhibit all of the above features. It has long been assumed that retrieval of information from long-term memory involves a parallel comparison of the retrieval cues in working memory to representations of information in long-term memory (e.g., Selfridge & Neisser, 1960). This parallel matching process can be modeled by constraint satisfaction, with constraints based on interconnections among features related to the set of retrieval cues (e.g., McClelland & Rumelhart, 1985). The use of a target analog as a retrieval cue for possible source analogs stored in memory is presumably a relatively complex special case of memory retrieval in general. The process of mapping two analogs (or an analog and a more abstract schema) can also be viewed as a process of constraint satisfaction. Our ACME model of mapping and ARCS model of retrieval (as well as the ECHO model of explanatory coherence) implement a constraint-satisfaction approach.

2.2 Constraints on Analogical Mapping

We will begin by considering the process of mapping two analogs to each other, after the source analog has either been spontaneously retrieved or presented by a teacher. Mapping is in some respects simpler than analog retrieval, as it involves only the target and one source, rather than the target and an indefinitely large pool of possible analogs stored in memory. It is certainly the case that people can often readily map a target analog to a source problem that they cannot easily retrieve (Gick & Holyoak, 1980, 1983; Holyoak & Koh, 1987; Keane, 1988). As we will see, however, the constraints that govern analogical mapping provide valuable guidance in developing a broader model that includes the retrieval step.

The core of any constraint-satisfaction theory is the specification of constraints. Three classes of constraints recur in theoretical treatments of analogy: *structural*, *semantic*, and *pragmatic*.

2.2.1 Structural consistency

Many theorists, particularly Gentner (1983), have stressed the importance of consistent structural correspondences as a criterion for an intuitively satisfying analogical mapping (Burstein, 1986; Falkenhainer, Forbus, & Gentner, 1986; Gick & Holyoak, 1980; Hofstadter, 1984; Winston, 1980). Studies have shown that greater structural consistency leads to greater ease of mapping (Gick & Holyoak, 1980; Holyoak & Koh, 1987; Ratterman & Gentner, 1987). Loosely speaking, a source analog can serve as a model for the target if objects in the two analogs can be placed into correspondence so that relations also correspond. A formal definition of structural consistency can be developed in terms of the concept of a *morphism* (Palmer, 1989). If we represent analogs as sets of interrelated propositions (Gentner, 1983), then structural consistency requires that if a proposition P in the target is in correspondence with a proposition P′ in the source, then the predicate and argument(s) of P must each correspond to the respective predicate and argument(s) of P′. Two analogs constitute an isomorphism if the mapping between them is structurally consistent and one-to-one. Importantly, it is not generally possible to decide whether any pair P and P′ are structurally consistent without considering the entire set of correspondences between propositions in the target and source. This interdependence inherent in the constraint of structural consistency is illustrates the need for parallel constraint satisfaction.

It is also important to realize that the kinds of analogies of psychological interest virtually never have the structure of a strict isomorphism. Rather, some elements of the target may have no apparent corresponding element in the source (or vice versa); some correspondences may be many-to-one (a homomorphism) or one-to-many (violating the formal definition of a function); and structural consistency may occasionally be violated. Nonetheless, useful naturalistic analogies intuitively can be viewed as *approximations* to isomorphisms (Holland et al., 1986). Constraint satisfaction provides a mechanism for treating structural and other constraints as continuous "pressures" (Hofstadter, 1984), rather than rigid restrictions, which is essential for finding imperfect but useful mappings.

2.2.2 Semantic similarity

Various theorists have suggested, and empirical evidence confirms, that object and predicate similarity influence the mapping process, with high semantic similarity leading to greater ease of mapping (Gentner & Toupin, 1986; Holyoak & Koh, 1987; Ross, 1987; Winston, 1980). Empirical evidence indicates that semantic similarity and structural consistency have distinct effects on the use of analogy. Semantic similarity has a more pronounced effect on the retrieval of a source analog than on the mapping process (Gentner & Landers, 1985; Holyoak & Koh, 1987; Ratterman & Gentner, 1987). In addition, although judgments of the aptness or soundness of analogies and metaphors are positively correlated with structural consistency, they are negatively correlated with similarity (Tourangeau & Sternberg, 1982). These separable effects of structural consistency and semantic similarity motivate treating the two kinds of constraints as distinct.

2.2.3 Pragmatic centrality

Another major type of constraint on mapping that many theorists have proposed involves the pragmatic importance of the elements of the two analogs—the assessment of relevance to the goals of the analogist (Holyoak, 1985). In general, people develop a mapping between two analogs in order to achieve some purpose, such as solving a problem, answering a question, or arguing for a desired conclusion. Some theorists have emphasized the centrality of causal knowledge in determining the most appropriate mapping (Winston, 1980); others have focused on the roles of high-level plans, goals, and functional knowledge (Burstein, 1986; Carbonell, 1983, 1986; Kedar-Kabelli, 1985). These proposals assume that the analogist uses explicit or implicit knowledge about the purpose the analogy is intended to serve to help direct the mapping process.

Although few would dispute that pragmatic knowledge influences the use of analogy, there remains disagreement as to the locus of its influence. Clearly pragmatic considerations weigh heavily in the initial selection of a plausibly useful source analog and in the subsequent transfer process. It is less clear that pragmatic constraints operate directly in the mapping process, as pragmatic knowledge is very difficult to separate empirically from structural and semantic information (Gentner, 1989). However, pragmatic guidance would increase the likelihood that inferences generated on the basis of an analogical mapping will in fact be relevant to the analogist's goals.

2.3 Analogical Mapping in ACME

The three basic constraints we have described—isomorphism (i.e., structural consistency and one-to-one mapping), semantic similarity, and pragmatic centrality—form the theoretical basis for the ACME model (Holyoak & Thagard, 1989b). ACME operates on propositional representations of the sort illustrated in Table 1, which represent two analogs used in experiments by Holyoak and Koh (1987). The "radiation problem" involves using rays to destroy a tumor without destroying surrounding healthy tissue, and the "lightbulb problem" requires using a laser to fuse a filament in a lightbulb without breaking the surrounding bulb. Each proposition consists of a predicate, its arguments, and an identifier for the proposition. Given such inputs, ACME builds a network of "mapping units" representing possible correspondences between elements (e.g., *ray-source=laser*). The units formed are restricted such that only propositions within corresponding major parts of the analogs (e.g., starting conditions and goals of problems) are mapped, and only elements of the same logical type (propositions, predicates, and arguments) are mapped.

Once the mapping units are formed, links are set up to enforce the various constraints. To capture structural consistency, excitatory links are formed between each proposition-mapping unit and the corresponding predicate-mapping and argument-mapping units, and also among the latter. Two special units, which are fixed at maximum activation, are used to enforce semantic and pragmatic constraints. The special semantic unit has excitatory connections to all predicate-mapping units, with weights ranging from a minimum value representing no similarity to a maximum value representing identity. As illustrated in Table 1, ACME can be given intermediate similarity weights for predicates that are similar but not identical, such as *ray-source* and *laser.*

The special pragmatic unit (not used in this example) can give extra excitation to mappings involving any element stated to be especially "important", or to any mapping unit that is "presumed" to hold in advance. In addition to representing such pragmatic information by weights, ACME also has the capacity to represent various types of questions and to selectively favor mappings that could provide relevant answers.

Once the network has been established, it is allowed to settle to a stable asymptote, using Grossberg's (1978) activation-updating procedure. The network established for the example in Table 1 consists of 178 mapping units and 1565 symmetric links. The network reaches asymptote after 31 cycles of updating, producing the set of "best" mappings of predicates and objects presented in Table 2 (taken directly from the output of the ACME program). The values given are asymptotic activation levels (with a maximum value of 1). If more than one mapping for an element achieves a relatively high activation, the output gives the alternative possibilities. ACME has been run on numerous other examples, including simulations of empirical data on human analogical mapping (Gentner & Toupin, 1986; Holyoak & Koh, 1987). Many of these analogies are considerably less isomorphic than are the analogs in Table 1.

2.4 Analog Retrieval in ARCS

The ACME program takes as its inputs predicate-calculus representations of two analogs, together with numerical weights reflecting semantic similarity and pragmatic centrality. No attempt is made to model conceptual knowledge stored in long-term memory, or the computation of semantic similarity from more basic information. The issues of long-term memory organization and similarity computation must necessarily be addressed, however, in any theory of the retrieval of analogs. The ARCS program (Thagard et al., 1989) is an attempt to combine aspects of our earlier PI model of analog retrieval (Holyoak & Thagard, 1989a; Thagard, 1988a) with the constraint-satisfaction approach of ACME.

Table 1

Predicate-Calculus Representations of Radiation and Lightbulb Problems

```
RADIATION PROBLEM (target)

Start:  (ray-source (obj_ray) r1)
        (tissue (obj_tissue) r2)
        (tumor (obj_tumor) r3 )
        (surround (obj_tissue obj_tumor) r4)
        (outside (obj_ray obj_tissue) r5)
        (can-produce (obj_ray obj_rays_high) r6)
        (high-intensity (obj_rays_high) r7)
        (can-destroy (obj_rays_high obj_tumor) r8)
        (can-destroy (obj_rays_high obj_tissue) r9)
        (can-produce (obj_ray obj_rays_low) r10)
        (low-intensity (obj_rays_low) r11)
        (cannot-destroy (obj_rays_low obj_tumor) r12)
        (cannot-destroy (obj_rays_low obj_tissue) r13)

Goals:  (destroy (obj_ray obj_tumor) r21)
        (not-destroyed (obj_tissue) r22)


LIGHTBULB PROBLEM (source)

Start:  (laser (obj_laser) b1)
        (bulb (obj_bulb) b2)
        (filament (obj_filament) b3)
        (surround (obj_bulb obj_filament) b4)
        (outside (obj_laser obj_bulb) b5)
        (can-produce (obj_laser obj_beams_high) b6)
        (high-intensity (obj_beams_high) b7)
        (can-fuse (obj_beams_high obj_filament) b8)
        (can-destroy (obj_beams_high obj_bulb) b9)
        (can-produce (obj_laser obj_beams_low) b10)
        (low-intensity (obj_beams_low) b11)
        (cannot-fuse (obj_beams_low obj_filament) b12)
        (cannot-destroy (obj_beams_low obj_bulb) b13)

Goals:  (fuse (obj_laser obj_filament) b21)
        (not-destroyed (obj_bulb) b22)


SIMILARITY:  (similar ray-source laser .08)
```

Table 2

Asymptotic Activation Values of Best Mappings of Predicates
and Objects in Radiation Problem to Those in Lightbulb Problem

```
Network has settled by cycle 31.
Test: TEST1   Total times: 32
Mon Jul 18  4:31:12 P.M. PDT 1988
Radiation and lightbulb problems.
Units not yet reached asymptote: 0
Goodness of network:  4.81
Calculating the best mappings after 32 cycles.
Best mapping of RAY-SOURCE is LASER. 0.69
Best mapping of TISSUE is BULB. 0.59
Best mapping of TUMOR is FILAMENT. 0.59
Best mapping of SURROUND is SURROUND. 0.77
Best mapping of OUTSIDE is OUTSIDE. 0.77
Best mapping of CAN-PRODUCE is CAN-PRODUCE. 0.88
Best mapping of HIGH-INTENSITY is HIGH-INTENSITY. 0.71
Best mapping of CAN-DESTROY is CAN-DESTROY. 0.57
     Mapping with CAN-FUSE is also possible:  0.40
Best mapping of LOW-INTENSITY is LOW-INTENSITY. 0.71
Best mapping of CANNOT-DESTROY is CANNOT-DESTROY. 0.57
     Mapping with CANNOT-FUSE is also possible:  0.40
Best mapping of DESTROY is FUSE. 0.71
Best mapping of NOT-DESTROYED is NOT-DESTROYED. 0.71
Best mapping of OBJ_RAYS_LOW is OBJ_BEAMS_LOW. 0.89
Best mapping of OBJ_RAYS_HIGH is OBJ_BEAMS_HIGH. 0.89
Best mapping of OBJ_TUMOR is OBJ_FILAMENT. 0.90
Best mapping of OBJ_TISSUE is OBJ_BULB. 0.90
Best mapping of OBJ_RAY is OBJ_LASER. 0.90
```

Empirical evidence suggests that the three basic types of constraints—structural, semantic, and pragmatic—that guide analogical mapping operate in the retrieval process as well. Their relative importance, however, differs across retrieval and mapping. In particular, semantic similarity has a relatively greater influence on initial retrieval of analogs than on mapping (Gentner & Landers, 1986; Holyoak & Koh, 1987; Ratterman & Gentner, 1987; Ross, 1987). The prominence of semantic constraints in determining analogical access is consistent with the ubiquitous role of semantic features as retrieval cues for information stored in long-term memory. Unless there is some minimal semantic overlap between the concepts in a target analog and those in some potential source stored in memory, there will likely be no retrieval pathways linking them. Without such retrieval paths, no amount of pragmatic relevance or structural correspondence will suffice to access the stored analog, even if a useful mapping could in fact be derived using non-semantic constraints.

ARCS, like PI, incorporates a conceptual network of the sort often postulated in models of human semantic memory. We assume that long-term memory contains representations of structured situations, such as problems or stories, which the program models with the kind of predicate-calculus representations exemplified in Table 1. In addition, the predicates in these structures are assumed to be tokens of concept types (e.g., the *laser* predicate in the lightbulb problem is a token of the long-term memory concept "laser"). Concepts are represented as frame-like structures with pointers to related concepts of various standard sorts (e.g., synonyms, superordinates, subordinates, and antonyms). The similarity of any two concepts can be computed as a function of their overlap in the network. The ARCS program bases its similarity estimates on information drawn primarily from an automated thesaurus, WordNet (Miller, Fellbaum, Kegl, & Miller, 1987). Although WordNet presumably provides only an approximation to human conceptual organization, it offers the methodological advantage of being derived by techniques motivated by very different and independent considerations than our application to analog retrieval. The semantic data base used in ARCS draws its organization in terms of synonyms, antonyms, superordinates, subordinates, and part-whole relations from WordNet; since WordNet is still under development, however, only about 70% of the semantic entries in ARCS are directly derived from WordNet.

The operation of ARCS involves two major steps. Retrieval begins when the target analog and its constituent predicates are activated. An initial search process finds pathways from the predicates in the target through intervening concept nodes (which may include multi-step links) to predicates in various structures stored in long-term memory. ARCS thus resembles the PI program in that it uses converging concept-based retrieval pathways to make an initial estimate of the likelihood that a stored structure is a useful analog of the target. ARCS goes beyond PI, however, in introducing a second major step in which structural and pragmatic constraints are integrated with information based on semantic similarity. In particular, it is possible for two structures to have many similar predicates, yet radically different patterns of predicate-argument structure, making the structures mere "clang associates" rather than useful analogs. PI (like most other previous retrieval models) is insensitive to structural constraints prior to retrieval. In contrast, human retrieval of analogs, although strongly influenced by semantic overlap, is also sensitive to consistent predicate-argument structure (Holyoak & Koh, 1987; Ratterman & Gentner, 1987).

Accordingly, ARCs has a second major step in which it builds a mapping network similar to that constructed by ACME. Specifically, for each proposition in the target with a predicate sufficiently similar to one in some stored structure, an excitatory subnetwork similar to those in ACME is formed linking the predicates, the propositions, and the corresponding arguments. In addition, a unit is created to represent the possible mapping of the target to the stored structure that includes the similar predicate. This structure-mapping unit is linked to the corresponding proposition-mapping unit. Such subnetworks are formed based on all the semantic links between the target and alternative stored structures. As in ACME, inhibitory links are established between inconsistent mapping units, and the network is allowed to settle. The alternative structures in memory thus compete to map onto the target, and the structure-mapping unit that first achieves sufficiently high activation becomes the "winner", an outcome that we assume corresponds psychologically to conscious access to a stored source analog. At this point the source is available for further processing by an ACME-like mapper.

In modeling the process of analog retrieval, ARCS introduces a number of features that seem more psychologically plausible than those embodied in ACME. First, it introduces a conceptual network that allows the program to compute inter-concept similarity from more elementary information. Second, it replaces ACME's purely syntactic restrictions on the formation of mapping units with semantic restrictions. ARCS typically forms far fewer units and links between the target and a given source analog than would ACME. Finally, ARCS captures one of the main properties of the earlier PI model, in that the output of the retrieval process provides the beginning of a complete mapping between the target and chosen source. That is, in the process of retrieving the source, correspondences based on similar predicates are already computed, so that further post-access mapping can begin with part of the mapping already established.

The ARCS program has been tested on data bases that include synposes of 25 Shakespearian plays (cf. Winston, 1980), and 100 of Aesop's fables, as well as materials that have been used in experiments on human analog retrieval (Holyoak & Koh, 1987; Ratterman & Gentner, 1987). In the domain of Shakespearian synopses, when ARCS is probed with a represen tation of West Side Story, it succeeds in retrieving the analogous play Romeo and Juliet.

3 IMPLICATIONS FOR THE ARMY

The work accomplished in this project has a number of significant implications for the development of new procedures for training and practice in the areas of military problem solving and decision making.

3.1 Training Implications

If our theories of analogical mapping and retrieval are correct, then analogical problem solving (case-based reasoning) can potentially be improved by attention to semantic, structural, and pragmatic constraints. Teaching is most directly a matter of mapping rather than retrieval, in that an instructor can explicitly point out a relevant analog to a student. Since, on our view, structural consistency is the most important component of analogical mapping, a teacher would be best advised to make sure that the two analogs are as isomorphic as possible, maintaining systematic relations between the phenomenon to be understood and the more familiar phenomenon that serves as its analog. But semantic and pragmatic matters should not be neglected. An instructor should try to pick an analog that has terms whose meanings are as similar as possible to those in the domain you are trying to explain. Most crucially, the teacher should avoid using semantically similar concepts that play *dis*analogous roles in the analogs. Attending to pragmatic constraints, the instructor should make clear the purpose of the analogy and how features of the source analog serve to explain or solve features of the target.

The usefulness of an analogy should, however, go beyond its immediate use in a specific training episode. The student should be able to recall the analogy later and be able to use it. Our theory of retrieval suggests that semantic similarity is the central constraint that guides retrieval, so instructors using analogs that they hope will be remembered later should make sure that the pragmatically important components of the two analogs are semantically similar to each other. In addition, based on the results of Catrambone and Holyoak (1988, 1989), it should be helpful to teach trainees to analyze multiple examples of a problem type in terms of specific, goal-oriented questions. Such procedures will foster the induction of problem schemas that are likely to be accessed even after a delay or context change.

3.2 Potential Aid for Computer-Assisted Problem Solving

The computational work accomplished in this project has provided the basis for the development of potential aids for decision making and problem solving. It may be possible to construct a computational device that uses analog retrieval and mapping, together with an evaluation of relative explanatory coherence, to assist a human problem solver. There is clear potential for use of such systems by army personnel to enhance their performance.

4 REFERENCES

Burstein, M. H. (1986). A model of learning by incremental analogical reasoning and debugging. In R. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artifical intelligence approach*, Vol. 2. Los Altos, Calif.: Kaufmann.

Carbonell, J. G. (1983). Learning by analogy: Formulating and generalizing plans from past experience. In R. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach.* Palo Alto, Calif.: Tioga Press.

Carbonell, J. G. (1986). Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In R. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artifical intelligence approach*, Vol. 2. Los Altos, Calif.: Kaufmann.

Catrambone, R., & Holyoak, K. J. (1988). Long-term analogical transfer. Paper presented at the meeting of the Psychonomic Society.

Catrambone, R., & Holyoak, K. J. (1989). Overcoming contextual limitations on problem-solving transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 1147-1156.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1986). The structure-mapping engine. *Proceedings of the Fifth National Conference on Artificial Intelligence.* Los Altos, Calif.: Morgan Kaufman.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science, 7*, 155-170.

Gentner, D. (1989). The mechanisms of analogical reasoning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning.* London: Cambridge University Press.

Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. In *Proceedings of the International Conference on Systems, Man and Cybernetics.* Tucson, AZ.

Gentner, D., & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science, 10*, 277-300.

Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12*, 306-355.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15*, 1-38.

Grossberg, S. (1978). A theory of visual coding, memory, and development. In E. L. J. Leeuwenberg & H. F. J. Buffart (Eds.), *Formal theories of visual perception.* New York: Wiley.

Hofstadter, D. R. (1984). The Copycat project: An experiment in nondeterministic and creative analogies. Cambridge, Mass.: MIT A.I. Laboratory Memo 755.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. (1986). *Induction: Processes of inference, learning, and discovery.* Cambridge, Mass.: Bradford Books/MIT Press.

Holyoak, K. J. (1985). The pragmatics of analogical transfer. In G. H. Bower (Ed.), *The psychology of learning and motivation*, Vol. 19. New York: Academic Press.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition, 15*, 332-340.

Holyoak, K. J., & Thagard, P. (1989a). A computational model of analogical problem solving. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning.* Cambridge, Mass.: Cambridge University Press.

Holyoak, K. J., & Thagard, P. (1989b). Analogical mapping by constraint satisfaction. *Cognitive Science, 13*, 295-355.

Holyoak, K. J., & Thagard, P. (1990). A constraint-satisfaction approach to analogue retrieval and mapping. In K. Gilhooly, M. Keane, R. Logie, & G. Erdos (Eds.), *Lines of thinking.* Chichester, U.K.: Wiley.

Keane, M. (1988). *Analogical problem solving.* Chichester, England: Ellis Horwood.

Kedar-Cabelli, S. (1985). Purpose-directed analogy. *Proceedings of the Seventh Annual Conference of the Cognitive Science Society.*

Maier, N. R. F. (1930). Reasoning in humans. I. On direction. *Journal of Comparative Psychology, 10*, 115-143.

Miller, G. A., Fellbaum, C, Kegl, J, & Miller, K. (1987). WordNet: An electronic lexical reference system based on theories of lexical memory. Princeton University Cognitive Science Laboratory Technical Report 11.

McClelland, J. L., & Rumelhart, D. E. (1985). Distributed memory and the representation of general and specific information. *Journal of Experimental Psychology: General, 114*, 159-188.

Palmer, S. E. (1989). Levels of description in information processing theories of analogy. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning.* London: Cambridge University Press.

Ratterman, M., & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society.* Hillsdale, N.J.: Erlbaum.

Ross, B. H. (1987). This is like that: The use of earlier problems and the separation of similarity effects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 629-639.

Schank, R. C. (1982). *Dynamic memory.* Cambridge, Mass.: Cambridge University Press.

Selfridge, O. G., & Neisser, U. (1960). Pattern recognition by machine. *Scientific American, 203*, 60-68.

Thagard, P. (1986). Parallel computation and the mind-body problem. *Cognitive Science, 10*, 301-318.

Thagard, P. (1988a). *Computational philosophy of science.* Cambridge, Mass.: Bradford Books/ MIT Press.

Thagard, P. (1988b). Dimensions of analogy. In D. Helman (Ed.), *Analogical reasoning.* Dordrecht, Holland: Kluwer.

Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences, 12*, 435-467.

Thagard, P., Cohen, D., & Holyoak, K. J. (1989). Chemical analogies. In *Proceedings of the Eleventh International Joint Conference on Articficial Intelligence.* San Mateo, CA: Kaufmann.

Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D. (1989). Analog retrieval by constraint satisfaction. Manuscript in preparation, Cognitive Science Laboratory, Princeton University.

Thagard, P., & Nowak, G. (1988). The explanatory coherence of continental drift. In A. Fine & J. Leplin (Eds.), *PSA 1988* (Vol. 1, pp. 118-126). East Lansing, MI: Philosophy of Science Association.

Tourangeau, R., & Sternberg, R. J. (1982). Understanding and appreciating metaphors. *Cognition, 11*, 203-204.

Winston, P. H. (1980). Learning and reasoning by analogy. *Communications of the ACM, 23*, 689-703.